



PATENT
Attorney Docket No.: 16869P-111500US
Client Ref. No.: 340301258US01

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:

Shinichiro YAMASHITA

Application No.: 10/806,986

Filed: March 22, 2004

For: STORAGE SYSTEM, CONTROL
METHOD FOR STORAGE
SYSTEM, AND STORAGE
CONTROL UNIT

Customer No.: 20350

Examiner: Unassigned

Technology Center/Art Unit: 3762

Confirmation No.: 2191

**PETITION TO MAKE SPECIAL FOR
NEW APPLICATION UNDER M.P.E.P.
§ 708.02, VIII & 37 C.F.R. § 1.102(d)**

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is a petition to make special the above-identified application under MPEP § 708.02, VIII & 37 C.F.R. § 1.102(d). The application has not received any examination by an Examiner.

(a) The Commissioner is authorized to charge the petition fee of \$130 under 37 C.F.R. § 1.17(i) and any other fees associated with this paper to Deposit Account 20-1430.

01/31/2005 AWONDAF1 00000040 201430 10806986

01 FC:1464 130.00 DA

(b) All the claims are believed to be directed to a single invention. If the Office determines that all the claims presented are not obviously directed to a single invention, then Applicants will make an election without traverse as a prerequisite to the grant of special status.

(c) Pre-examination searches were made of U.S. issued patents, including a classification search and a key word search. The classification search was conducted on or around September 9, 2004 covering Class 707 (subclasses 202 and 204), Class 711 (subclasses 114 and 162), and Class 714 (subclasses 6 and 15), by a professional search firm, Lacasse & Associates, LLC. The key word search was performed on the USPTO full-text database including published U.S. patent applications. The inventors further provided three references considered most closely related to the subject matter of the present application (see references #6-8 below), which were cited in the Information Disclosure Statement filed with the application on March 22, 2004.

(d) The following references, copies of which are attached herewith, are deemed most closely related to the subject matter encompassed by the claims:

- (1) U.S. Patent No. 6,629,264 B1;
- (2) U.S. Patent No. 6,732,243 B2;
- (3) U.S. Patent Publication No. 2003/0187947 A1;
- (4) U.S. Patent Publication No. 2004/0078644 A1;
- (5) PCT Patent Publication No. WO 2004/051479 A2;
- (6) U.S. Patent No. 6,601,187;
- (7) U.S. Patent Publication No. 2002/0095489 A1; and
- (8) U.S. Patent Publication No. 2003/0033523 A1.

(e) Set forth below is a detailed discussion of references which points out with particularity how the claimed subject matter is distinguishable over the references.

A. Claimed Embodiments of the Present Invention

The claimed embodiments relate to replication technology for storing the replication of data, which is stored in a storage volume of a storage unit of a computer system at a main site, to a storage volume in a storage unit of a storage system at a remote site.

Independent claim 1 recites a storage system comprising a first storage unit having a first storage volume for storing data; and a second storage unit communicably coupled to the first storage unit and having a second storage volume for storing data. The first storage unit includes a data transmission unit configured to transmit replicated data to a storage unit when data is written to the first storage volume. The second storage unit further includes a data reception unit configured to receive the replicated data and writing the replicated data to the second storage volume. The first storage unit further includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time. The second storage unit further includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit.

Independent claim 9 recites a method for controlling a storage system which system includes a first storage unit having a first storage volume for storing data, and a second storage unit in communication with the first storage unit and having a second storage volume for storing data, wherein the first storage unit includes a data transmission unit configured to transmit replicated data to the second storage unit when the data is written to a first storage volume, and the second storage unit includes a data reception unit configured to receive the replicated data and writing the replicated data to the second storage volume. The method comprises in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume.

Independent claim 17 recites a storage system comprising a first computer system including a first storage unit having a first storage volume for storing data, and a first information processing unit communicably coupled to the first storage unit; and a second computer system including a second storage unit having a second storage volume for storing data, and a second information processing unit communicably coupled to the first storage

unit. The first storage unit includes a data transmission unit configured to transmit replicated data to the second storage unit when the data is written to the first storage volume. The second storage unit includes a data reception unit configured to receive the replicated data and writing the replicated data to the second storage volume. The first storage unit includes a disk heart beat creation unit configured to repeatedly create a first heart beat message, and a disk heart beat write unit configured to repeatedly write the first heart beat message to the first storage volume at intervals. The second storage unit further includes a disk heart beat detection unit configured to detect the replicated first heart beat message, and a disk heart beat detection result transmission unit configured to transmit a signal indicating receipt of the replicated first heart beat message by the disk heart beat detection unit to the second information processing unit. The first information processing includes a node heart beat creation unit configured to repeatedly create a second heart beat message, and a node heart beat write request unit configured to repeatedly transmit a request to write the second heart beat message to the first storage volume. The first storage unit includes a node heart beat write unit configured to write the second heart beat message to the first storage volume according to the write request of the second heart beat message. The second storage unit includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit. The second information processing unit includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system.

Independent claim 19 recites a first storage control unit communicably coupled to a second storage control unit for controlling configured to control reading and writing of data to first, second and third storage volumes. The first storage control unit includes a data transmission unit configured to transmit replicated data to the second storage control unit when data is written to the first storage volume; a disk heart beat write unit configured to write a first heart beat message to the first storage volume; a data reception unit

configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume.

One of the benefits that may be derived is that it is possible to specify the failure area of the computer system more easily, in more detail, and more accurately.

B. Discussion of the References

1. U.S. Patent No. 6,629,264 B1

This reference discloses a controller-based remote copy system with logical unit grouping. Discussed is a remote copy set operation. A local host computer 101 requests a storage array I/O operation and a local array controller 301 presents a local volume that is part of the Remote Copy Set to the local host 101. The host 101 performs writes to the local volume on the local array 203, which copies the incoming write data to the remote volume on the target array 213. A flow diagram shows the operation of array controller 'heartbeat' timers. At step 600, in response to a write request, array controller A1 sends a write command and the host write data to target array controller B1 via fabric 103A so that the data is backed up on array 213. At step 605, controller A1 starts a command ('heartbeat') timer which keeps track of the time between issuance of the write command and a response from the target controller B1. If link 1 and controller B1 are operational, then controller B1 writes the data to array 213. During normal operation, at step 640, controllers C and C! periodically send pings to each other via dual asynchronous receiver/transmitters located at both ends of bus 330. See figure 3, 6A; column 8, lines 56-62; column 9, lines 43-51; and column 10, lines 45-48.

The reference merely discloses controllers that periodically send pings to each other via dual asynchronous receiver/transmitters located at both ends of the bus. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit,

repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

2. U.S. Patent No. 6,732,243 B2

This reference discloses data mirroring using shard buses. Discussed are controller management modules (CMMs). CMMs 104 and 108 mirror data to provide cache coherency to the network storage apparatus 100a. This can be accomplished by providing a DMA engine 188 in each CMM 104, 108 and a shared path 216 to send data to the other CMM 104, 108. A failover reset link 240 is also present between CMM 104 and CMM 108. Each CMM 104, 108 maintains a heartbeat signal which is communicated over the failover link 240, and monitored by the other CMM 104, 108. If a problem is detected in the heartbeat signal, a CMM 104, 108 can send a signal over the failover reset link 240 to terminate the operation of the other CMM 104, 108. See figures 2, 3; column 10, lines 3-5, 8, 10-12; and column 11, lines 35-40.

The reference merely discloses data mirroring using shared buses. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to

detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

3. U.S. Patent Publication No. 2003/0187947 A1

This reference discloses system and method for multi-destination merge in a storage area network. Disclosed is a data transfer protocol that includes a heartbeat MFC that is periodically transmitted from each controller 101 to each other controller 101. A controller 105 that receives a write request from a host 102 in the designated source controller for the received request. Once the entire write operation is received in the primary cache, it is copied to the mirror cache in operation 705. In operation 707, the mirror controller receives the operation into a receive buffer and in operation 709 the mirror cache receives the operation. After successfully completing other steps, in step 719, the write operation is drawn from the mirror cache into one or more transmit buffers. The data can be marked valid and the operation can be committed to media-based storage. See figures 7, 9; and paragraphs 59, 60, 65, and 69.

The reference merely discloses a data transfer protocol including a heartbeat MFC that is periodically transmitted from each controller to each other controller. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

4. U.S. Patent Publication No. 2004/0078644 A1

This reference discloses system and method for bi-directional failure detection of a site in a clustering system. Discussed is a heartbeat check 101 that uses application program interfaces when sending a heartbeat message. Heartbeat check 101a and 101b can be modules. Each host 100a has a clustering program 104a, a heartbeat check 101a, and an operating system 102a. Storage systems 110a and 110b are connected to each other by one or more remote links 150 so that communication can occur. When the heartbeat check 101a transmits a conventional change state command to the heartbeat volume 111a, the storage

system 110a changes the state of heartbeat volume 111a from a primary state to a storage state. Storage system 110a communicates the change in state to storage system 110b via remote link 150 so that storage system 110b can change the state of heartbeat volume 111b between a storage state and a primary state. When application 103a updates data on the user's PVOL 112a, the storage system 110a writes the data to the user's SVOL 112b by use of a conventional remote copy mechanism to transmit the data across remote link 151 to storage system 110b. See figure 1; and paragraphs 29, 31, 37, 38, 41.

The reference merely discloses a heartbeat check that uses application program interfaces when sending a heartbeat message. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

5. PCT Patent Publication No. WO 2004/051479 A2

This reference relates to a heartbeat mechanism for cluster systems. Discussed is heartbeat logic 835 that is programmed to generate and transmit a heartbeat message within a predetermined time interval. Heartbeat messages from each node are collected and stored in a quorum file 840. Each node in cluster 800 is allocated address space within the quorum file 840 to which its heartbeat messages are stored. The nodes access files on one or more data storage devices over a network. Files and or data are logically shared among the nodes with each database instance having access to all data. See abstract and paragraphs 41, 52, 54, 56; and figure 8.

The reference merely discloses a heartbeat mechanism for cluster systems. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit

configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

6. U.S. Patent No. 6,601,187

This reference discloses a data replication system having a redundant configuration including dual Fibre Channel fabric links interconnecting each of the components of two data storage sites, wherein each site comprises a host computer and associated data storage array, with redundant array controllers and adapters. Each array controller in the system is capable of performing all of the data replication functions, and each host 'sees' remote data as if it were local. Each array controller has a dedicated link via a fabric to a partner on the remote side of the long-distance link between fabric elements. Each dedicated link does not appear to any host as an available link to them for data access; however, it is visible to the partner array controllers involved in data replication operations. These links are managed by each partner array controller as if being 'clustered' with a reliable data link between them.

The reference merely discloses a data replication system in which each array controller has a dedicated link. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second

computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

7. U.S. Patent Publication No. 2002/0095489 A1

This reference discloses a cluster computing system, comprises: a production host group; a standby host group coupled to the production host group by a network; and a remote mirror coupled between the production host group and the standby host group, the remote mirror including a production site heartbeat storage volume (heartbeat PVOL) and a standby site heartbeat storage volume (heartbeat SVOL) coupled by a remote link to the heartbeat PVOL, with the production host group configured to selectively send a heartbeat signal to the standby host group by use of at least one of the network and the remote link. A method of checking for failure in a cluster computing system, comprises: generating a heartbeat signal from a production host group; selectively sending the heartbeat signal to the standby host group from the production host group by use of at least one of a network and a remote link; and enabling the standby host group to manage operations of the cluster computing system if an invalid heartbeat signal is received by the standby host group from the production host group.

The reference merely discloses a method of checking for failure in a cluster computing system that includes a remote mirror having a production site heartbeat storage volume and a standby site heartbeat storage volume coupled by a remote link. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that

includes a node heart beat transmission unit configured to transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

8. U.S. Patent Publication No. 2003/0033523 A1

This reference relates to system and method that improves security of a computer storage system by requiring an initiating computer to periodically reaffirm its identity by transmitting a message to a servicing computer. The message contains a previously established authentication message and a sequence value, established by and known only to the original participants. A message must be received by the servicing computer within a predetermined time interval in order to maintain data communications between the original participants.

The reference merely discloses a technique to improve security of a computer storage system. It does not teach a first storage unit that includes a disk heart beat write unit configured to repeatedly write a first heart beat message to the first storage volume at intervals within a predetermined time; and a second storage unit that includes a disk heart beat detection unit configured to detect a replication of the first heart beat message to be written to the second storage volume by the data reception unit, as recited in claim 1. Nor does it disclose in the first storage unit, repeatedly writing a first heart beat message to the first storage volume at intervals; and in the second storage unit, detecting the replicated first heart beat message to be written to the storage volume, as recited in claim 9. It further fails to teach a second storage unit that includes a node heart beat transmission unit configured to

transmit to the second information processing unit the replication of the second heart beat message written to the second storage volume by the data reception unit; and a second information processing unit that includes a node heart beat detection unit configured to detect the replication of the second heart beat message, and an operation status unit configured to determine operational status of the first computer system according to the second heart beat message and the first heart beat message, and a fail-over execution unit configured to transfer information processing from the first computer system to the second computer system according to the operational status of the first computer system, as recited in claim 17. It also fails to teach a data reception unit configured to receive a replicated second heart beat message written to the second storage volume by the second storage control unit, and in response write the second heart beat message to the third storage volume; and a disk heart beat detection unit configured to detect the second heart beat message written to the third storage volume, as recited in claim 19.

(f) In view of this petition, the Examiner is respectfully requested to issue a first Office Action at an early date.

Respectfully submitted,



Chun-Pok Leung
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
Attachments
RL:rl
60396308 v1